



An Integrated Instrumentation Architecture for NGI Applications

Dan Reed Ruth Aydt

University of Illinois

Ian Foster Darcy Quesnel Steven Tuecke

Argonne National Laboratory

<http://www-pablo.cs.uiuc.edu/Project/Pablo/NGIOverview.htm>



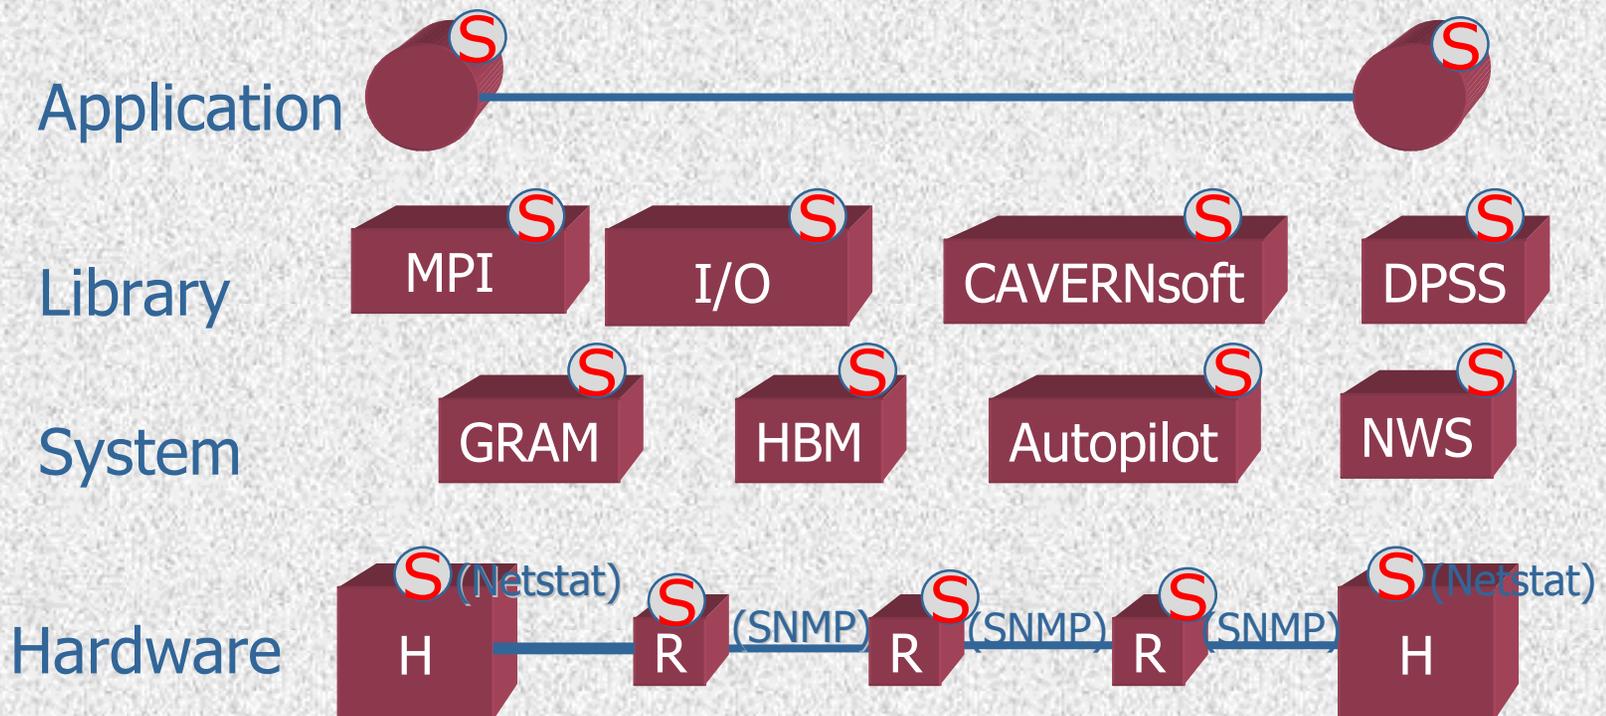
Project Goals

- Produce uniform mechanisms
 - instrumentation and event notification
 - » qualitative and quantitative data
 - dynamic adaptation
- Catalyze development of both network-aware middleware and sophisticated network-aware applications

“A Uniform Instrumentation, Event, and Adaptation Framework for Network-Aware Middleware and Advanced Network Applications”

Multilevel Sensor Example

- Multilevel data needed for analysis
 - possible performance problems at all levels
 - diverse data sources; *no standard access mechanisms*
 - *no standard publication or discovery techniques*

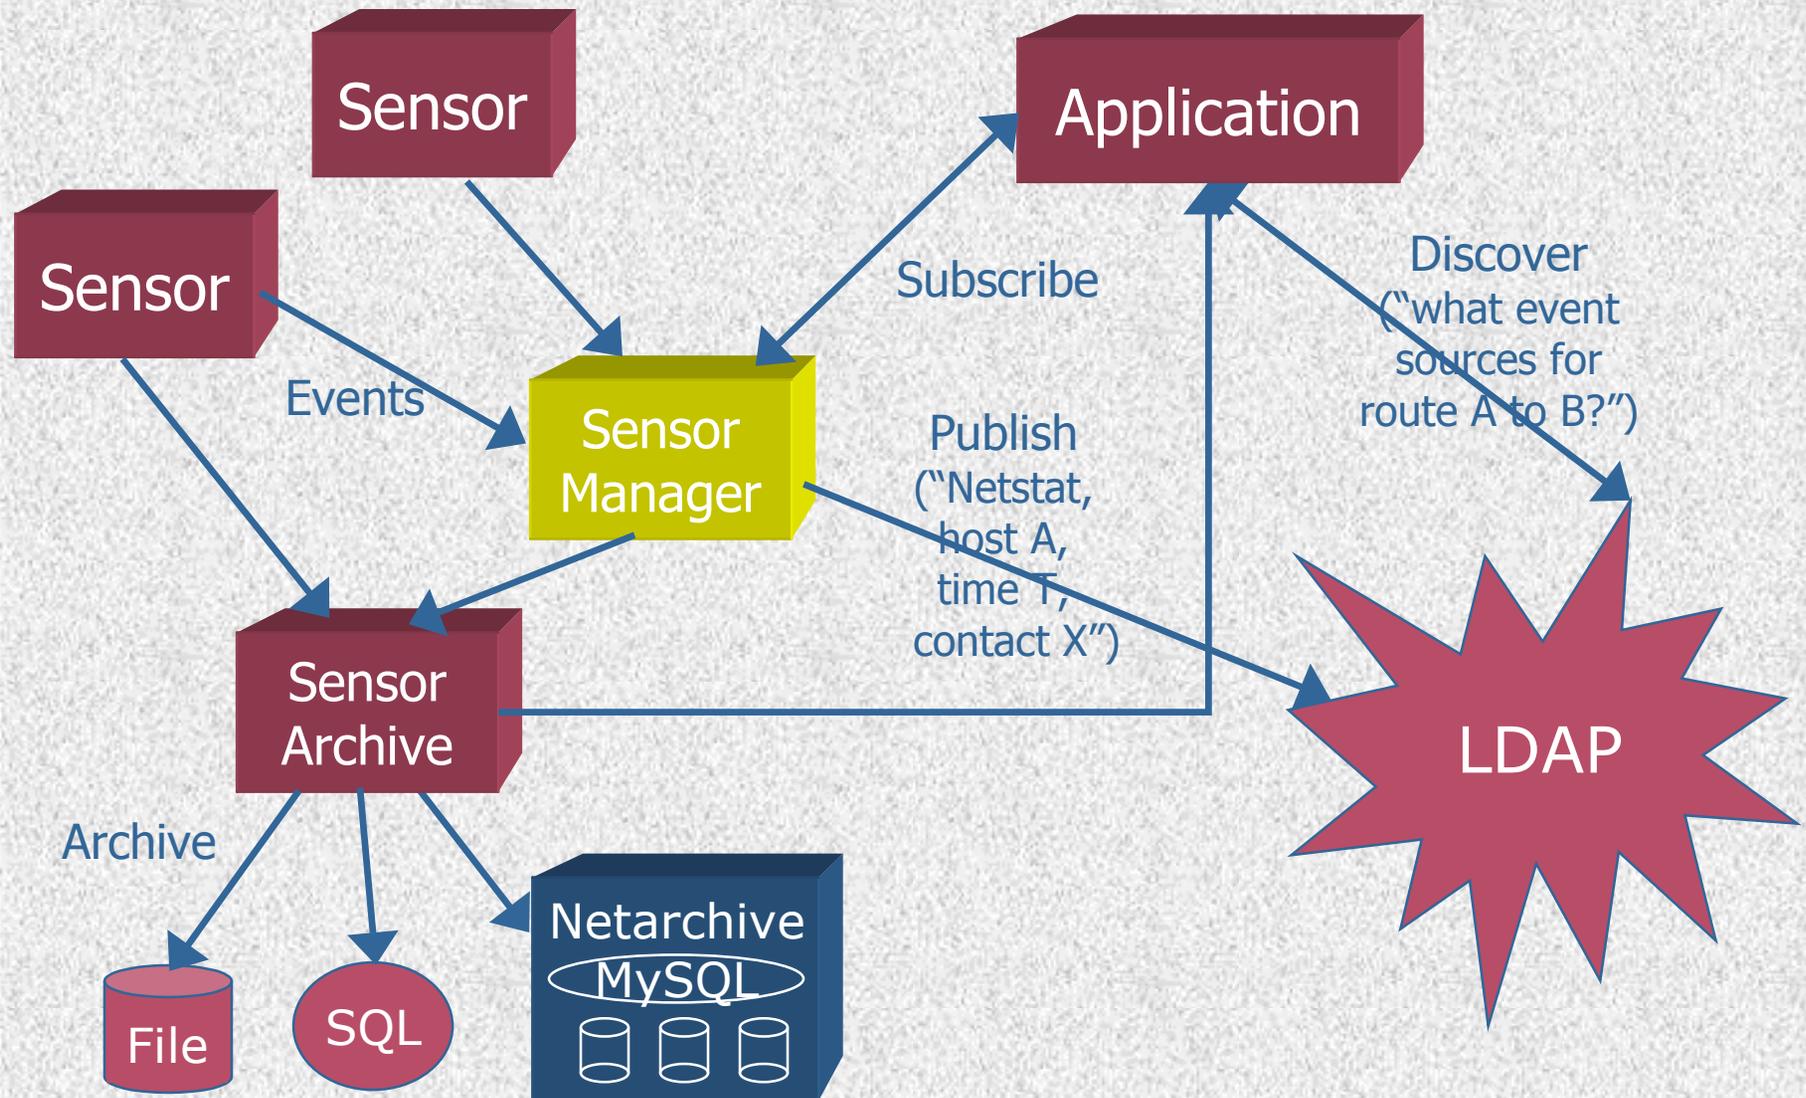




Key Technical Innovations

- **Sensor mechanisms**
 - creation, publication, discovery and access
- **Synthesis and analysis techniques**
 - extraction of qualitative behavior and trends
- **Adaptation techniques**
 - exploitation of sensor data
 - optimization of middleware and applications
- **Implementation mechanism**
 - *Globus/Autopilot/Netlogger integration*

Instrumentation Architecture





Project Sensor Approach

- **Directory service (LDAP)**
 - publish: **source, type, contact [online, archive]**
 - discover: **"find all event sources of type X"**
- **Autopilot sensor manager extensions**
 - publication, subscription, and archiving
- **Standard data formats**
 - LBL Netlogger, Illinois SDDF, and XML
 - » **standard converters (e.g., SDDF to XML, Netlogger to SDDF)**
- **Relational database archive**
 - publicly available SQL implementation
- **Standard sensor set integration**



Sensor Publication and Discovery

■ Globus LDAP MDS

- Metacomputing Directory Service (MDS)
- scalable, global infrastructure for publishing and discovering sensor managers

■ Approach

- sensors send attributes to sensor manager
- sensor manager publishes availability via LDAP
- clients discover sensor managers from LDAP
 - » then directly subscribe to current or archived sensor data

■ Netlogd/Globus/Autopilot extension/integration



Archiving Sensor Streams

■ SQL database

- each event as a record in an SQL database
 - » offers rich query support

■ Netarchive

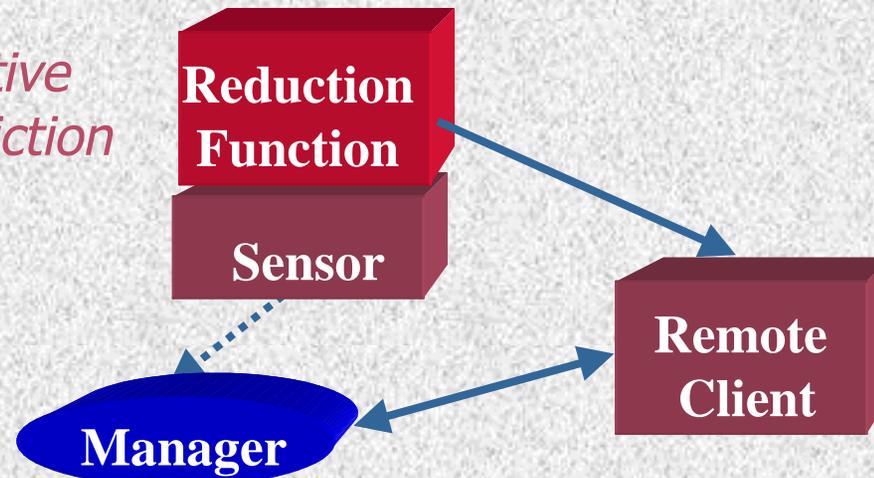
- each event stored in a file with SQL index
 - » offers performance and scale

■ We will explore SQL databases

- emphasize sensor data reduction at sources
- reduce event data volume for archiving
- prototype XML to SQL interface operational

Standard Sensors: Autopilot Base

Quantitative and qualitative data reduction and prediction



■ Quantitative sensors

- application
 - » software and hardware
- library
 - » MPI, I/O, HDF, and MPI-IO
- daemon
 - » network & system statistics

■ Software

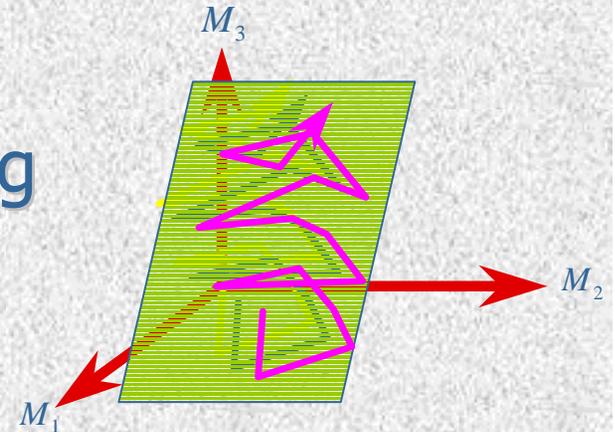
» Netlogger, Globus, Autopilot, ...

■ Two aspects

- quantitative resource use
 - » numerical measurements
- qualitative request patterns
 - » behavioral classification

Data Reduction Techniques

- Challenge: reduce sensor data volume
 - many metrics and concurrent activities
- Statistical clustering
 - based on square error clustering
 - reduces the number of points
- Projection pursuit
 - based on principal component analysis
 - identifies “important” metrics
- Result
 - relevant metrics from relevant sites

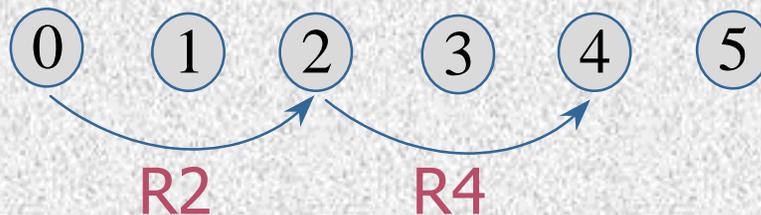




Classification and Prediction

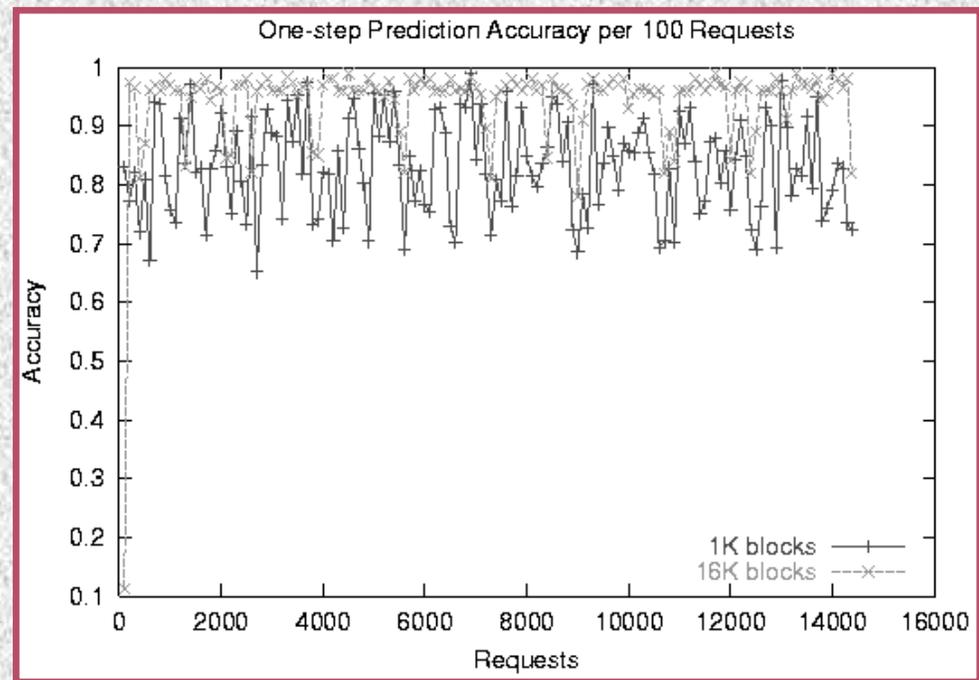
- Two axes for classification and prediction
 - spatial (where) and temporal (when)
- Neural network classification (ANNs)
 - accepts quantitative sensor data
 - generates qualitative classification
 - » *regular, irregular, large, small, bursty, slow, fast ...*
- Hidden Markov models (HMMs)
 - learns access probability distribution functions
 - recognizes non-qualitative patterns
- ARIMA time series
 - learns temporal behavior and predicts future patterns

LLNL ALE3D HMM I/O Prediction

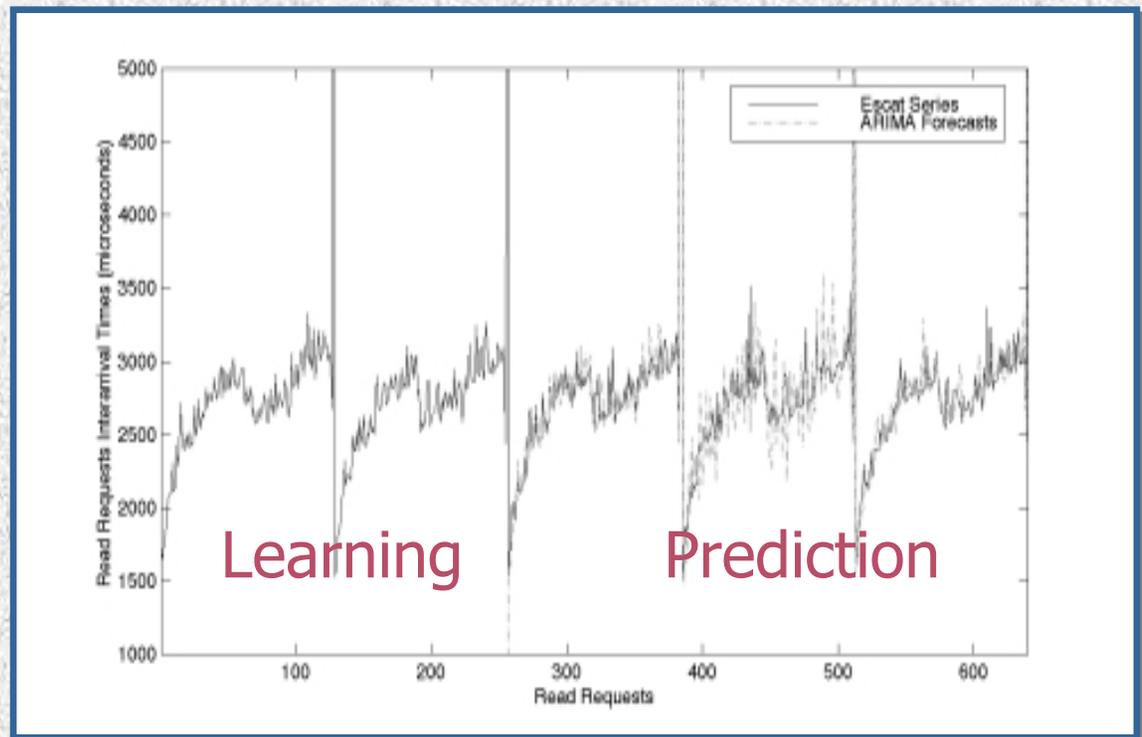
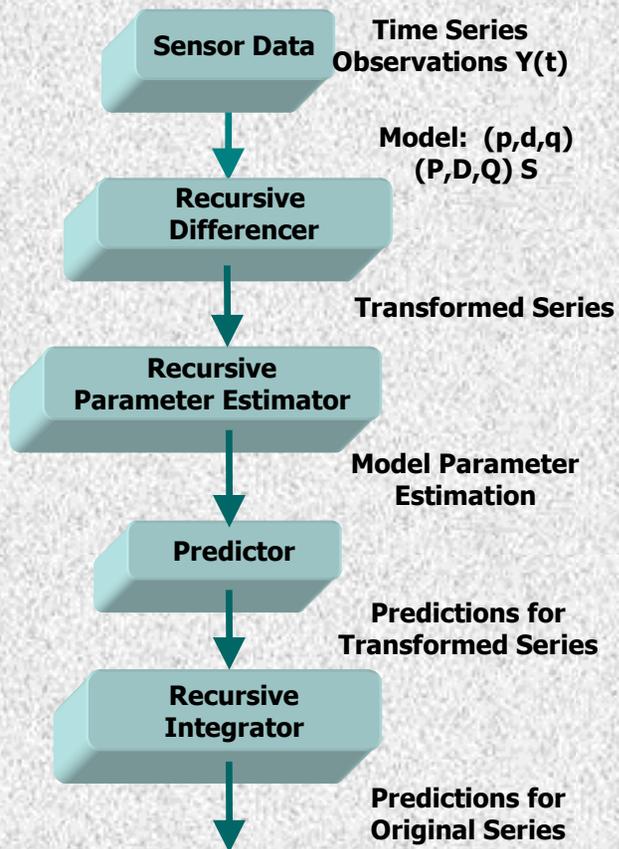


Example:
Read only strided

- Prediction
 - I/O block accesses
 - high accuracy
- Other domains
 - network traffic
 - system utilization

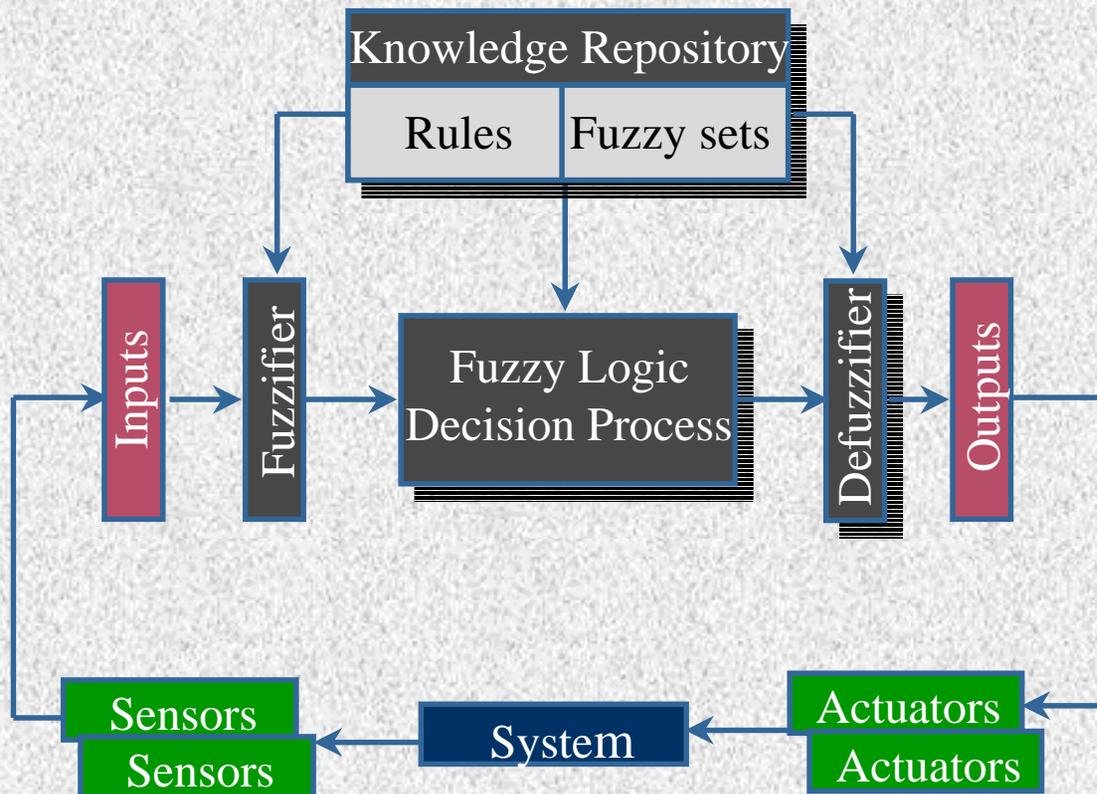


Caltech ESCAT ARIMA Predictions



Additional funding from DOE ASCI and NSF

Middleware Adaptation Process



■ Fuzzy rule base

- qualitative behaviors
- retargetable

■ Catastrophe theory

- rule optimization
- transitions
- hysteresis
- near-optimal control

■ *Result*

- *software control toolkit*

Based on Autopilot toolkit



Security

■ Grid Security Infrastructure (GSI)

– will be used throughout

- » “manager M accepts only streams from sensors of user U”
- » “manager N only publishes streams to clients of users A, B, C”

■ As a first step

– LBNL augmented Netlogger C client with GSI

Initial Applications

- Replica creation in “data grid” applications
 - online and historical instrumentation
 - » large data transfers (application, library, and network)
 - DPSS and Globus-IO (with LBNL)
 - application-level selection of replicas
 - » based on sensor information
- MPI video streaming (Karonis and Papka)

Project Timeline

